Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 16:01:35
REFERENCE NO: 273

Author Names & Affiliations

- Ellen Rathje - The University of Texas at Austin
- Dan Stanzione - The University of Texas at Austin

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Natural Hazards Engineering, Civil Engineering

Title of Submission

Future Needs for Advanced Cyberinfrastructure: A natural hazards engineering perspective from DesignSafe

Abstract (maximum ~200 words).

The DesignSafe project is the cyberinfrastructure component of the NSF Natural Hazards Engineering Research Infrastructure (NHERI) community. Researchers in this community seek to improve the design of buildings, roads, and bridges to minimize the impact of natural hazards on both property and human life. The NHERI community faces a variety of research challenges. As the scale and density of data collected from experiments and field observations increases, a more sophisticated approach to data management and analysis is needed that includes data discovery as a first-class object. And, in order to fully understand damage vectors and develop remediation approaches the NHERI community must develop a scientific approach that integrates phenomena across scales from hundreds of miles to inter-particle forces. Along the way, new computational approaches are needed to enable numerical simulations that will test, validate, and ultimately extend the usefulness of this new science. These challenges will only be met with a robust, modern, reliable cyberinfrastructure that includes sophisticated data management and discovery, new algorithms and software, and high performance computing. In order to have the most impact on our research agenda, the CI community must provision these services to meet the needs of users with different degrees of computational sophistication.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

The DesignSafe project is the cyberinfrastructure component of the NSF Natural Hazards Engineering Research Infrastructure (NHERI) community. This community is concerned with the resiliency of the built environment in the face of natural hazards presented by earthquakes, wind, and water. The community is large, consisting of a broad multidisciplinary set of engineers and scientists who perform controlled laboratory experiments, conduct simulations, and collect field observations of damage. The community also interacts closely with those in the geoscience community who seek to understand and forecast natural hazards. Researchers in the natural hazards engineering community seek to improve the design of buildings, roads, and bridges, and ultimately to enhance building codes, to minimize the risks and impacts of hazards on both property and human life.

NHERI includes, in addition to the DesignSafe cyberinfrastructure, a collection of large-scale experimental facilities, a rapid response reconnaissance facility, and a simulation center. The vision of the NHERI project is to facilitate transformative research that could not be achieved without the integration of the NHERI components, and the DesignSafe cyberinfrastructure plays a critical role in that integration. The research challenges in natural hazards engineering encompass a wide range of issues. The increasing scale and density of data collected from large-scale laboratory experiments and field investigations impedes data interpretation using traditional methods, and requires more sophisticated approaches for data management and analysis. These new approaches require access to large-scale data repositories that are searchable and provide the data in a well-documented and well-organized manner, as well as high performance computing resources to perform the transformations and exploration necessary to yield the most insight from these data.
Numerical simulations of the response of infrastructure to natural hazards play a critical role in evaluating the resilience of the built environment. Enhancing the fidelity of numerical simulations to better capture the response of these complex physical systems will improve our ability to forecast the impact of natural hazards on the homes and businesses in the zone of impact. Improving numerical simulations requires robust, large-scale, and modern high performance computers to conduct the simulations, as well as access to high-quality datasets for validation.

Numerical simulations across multiple scales are required to better understand regional and system-level responses. For instance, in the case of earthquakes it is important to properly model the ground shaking at the regional scale, then couple these estimates of ground shaking with models of the response of an individual structure at site scale, and then to link the computed structural response to a model of the building materials at the nanoscale to predict damage or failure. This work requires complex interactions between various simulation codes, as well as with input data and material characterization parameters, to provide robust assessments of the regional-scale effects of natural hazards.

Responding in a timely fashion in the wake of an event is critical – to rapidly model the threat of aftershocks or other continuing threats, to assess the stability and safety of buildings and roads, and to mount a coordinated emergency response.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

In recent years, we have seen the needs for cyberinfrastructure in our disciplines grow in breadth, depth, and complexity. For those working in large scale simulation, the appetite for high-end computing cycles is voracious and growing seemingly exponentially. But, in addition to simulation needs, we also note breadth in demand for other cyberinfrastructure services as well. The use of computing now takes many forms for both simulation and data analysis – researchers in this field use HPC, high throughput computing, as well as interactive computing to explore data through Matlab, Labview and, increasingly, Jupyter notebooks. Related to interactive computing, there are increasing needs for defined quality of service – to generate results with a real-time turnaround constraint, perhaps even at the expense of fidelity or accuracy in the result. Each of these computing modes is distinct and has its own needs and user communities, and is suitable for a different (though sometimes overlapping) range of problems. Likewise, each of these modes may require new and different algorithms to properly use the rapidly evolving hardware that comprises these systems.

Data is also a significant issue. Like computation, the demand for data storage is also large and growing. Increasingly, demand goes beyond simple storage to a comprehensive data management system that supports a secure, collaborative, and reproducible environment for research data. Ideally, this data system is coupled to our computational platforms, allowing analytics to be performed on data sets, with provenance that makes these numerical experiments reproducible. Researchers wish to be able to share their data, but need a low-barrier way to share large amounts of data without extensive manual intervention. “Sharing” in this context means several different things, from sharing raw data and early results with collaborators and colleagues immediately, to ultimately sharing publication data publicly in a way that is documented, understandable, and (perhaps most importantly) discoverable. Search mechanisms to locate relevant data are important.

Once the data management and computation infrastructure is in place, visualization is also important to transform measured and computed data into knowledge. Increasingly, researchers are curious about the role machine learning may have in mining additional insights from data.
Finally, a simplified user interface, and the skills needed to make use of all of the above capabilities, are a must (see “other considerations” below). The systems and in particular the software tools used to provide these capabilities are changing rapidly, and are diverse; simply selecting the appropriate tool for a job is becoming a difficult task for many researchers. Efforts by NSF should supply not just capability, but predictable, stable capability wherever possible – when tools and systems, and their providers, change rapidly (on the scale of “shorter than my grant program”), it is tempting to develop in-house capability rather than use those NSF provides in the national cyberinfrastructure.

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

A significant issue we have identified is the skillset and workforce required to properly match available cyberinfrastructure capabilities to the problems we have at hand. Seldom is this as simple as the two most commonly applied extremes: “read the documentation” on one hand, and “hire a computer person” on the other. Neither solution is adequate.

Modern cyberinfrastructure systems can by dauntingly complex, with computers containing tens of thousands of cores, alternative accelerated architectures, and many kinds of storage to choose from. Software systems are equally complex, with simulation programs comprising hundreds of thousands of lines of code and dozens of configuration options. Other parts of the cyberinfrastructure ecosystem have equal complexity (including the options for many different providers of these services).

We observe an increasing divergence of researchers making use of CI into two camps – the “CI ninjas” who make use of high performance computing, parallel programming, and sophisticated numerical methods, and the “normal users” who work (frequently) in Matlab and other visual environments at the desktop scale. While one camp can make use of giant HPC systems with hundreds of thousands of cores and deep memory and storage hierarchies, these systems, and the programming models that go with them, are increasingly foreign to the other camp, who are focused primarily on ease of use. Simple training webinars and documentation will not bridge the gap.

Often, bringing in a “CI expert” is considered a solution in this case – however, this term is vague and not always well-defined, and those doing the hiring may not be aware of the skillset needed. Hiring a “big data” person to deal with parallel processing may result in Hadoop being chosen for a problem where a more traditional MPI program would be orders of magnitude more efficient. Computing is a broad field, and hiring a single programmer for a project limits you to the vision of the particular candidate. The old adage “when you only have a hammer, everything begins to look like a nail” often applies.

A better understanding of computational science and engineering as a discipline, and some indication that persons are trained in particular areas of it would be helpful – including investments in insuring these trained persons exist. Moreover, mechanisms that allow the sharing of expertise between these professionals are also vital: perhaps by building critical mass of the necessary kinds of expertise at centers that can be accessed by numerous researchers, or through virtual organizations that allow the “computer person” in a given project to not operate on an island.